

Multistage Adaptive Estimation of Sparse Signals

Dennis Wei and Alfred O. Hero, III

Abstract

This paper considers sequential adaptive estimation of sparse signals under a constraint on the total sensing effort. The advantage of adaptivity in this context is the ability to focus more resources on regions of space where signal components exist, thereby improving performance. A dynamic programming formulation is derived for the allocation of sensing effort to minimize the expected estimation loss. Based on the method of open-loop feedback control, allocation policies are then developed for a variety of loss functions. The policies are optimal in the two-stage case and improve monotonically thereafter with the number of stages. Numerical simulations show gains up to several dB as compared to recently proposed adaptive methods, and dramatic gains approaching the oracle limit compared to non-adaptive estimation. An application to radar imaging is also presented.

Index Terms

Adaptive sensing, adaptive sampling, resource allocation, sparse signals, dynamic programming.

I. INTRODUCTION

Adaptive sensing and inference have been gaining interest in recent years in signal processing and related fields. Potentially substantial gains in performance can be achieved when observations are made sequentially and adaptively, making use of information derived from previous observations. This work focuses on sparse signals, i.e., signals that occupy a small number of dimensions in an ambient space. It is now well-known that compressed sensing offers an efficient non-adaptive strategy for acquiring sparse signals, relying on a relatively small number of observations that are incoherent with the basis in which

This work was partially supported by Army Research Office grant W911NF-11-1-0391.

The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: dlwei@eecs.umich.edu, hero@eecs.umich.edu).

the signal is sparse (see e.g. [1], [2]). However, when noise is present and sensing resources are limited, incoherent observations may not be the most efficient since a large fraction of the resources are allocated to dimensions where the signal is absent. Alternatively, by allocating resources according to estimates of the signal support obtained from past observations, better signal-to-noise ratios (SNR) are possible. Applications in which adaptive sensing of sparse signals can be readily utilized include surveillance using active radars [3], [4], spectrum sensing in cognitive radio [5], [6], and gene association and expression studies [7].

Existing methods for adaptive sensing of sparse signals can be roughly grouped around two classes of models. In the first class, which is the focus of this paper, observations are restricted to single components in the basis that induces signal sparsity, while resources can be distributed arbitrarily over components and observation stages. An optimal two-stage allocation policy was developed in [3] for a cost function related to bounds on estimation and detection performance. Subsequent developments stemming from [3] include a modification to handle non-uniform signal priors [8], a simplification based on Lagrangian constraint relaxation [9], and a multiscale approach that uses linear combinations in the first stage to reduce the number of measurements [4]. Based on a similar model but in a different direction, a method known as distilled sensing [10] was proposed for signal support identification and was shown to be asymptotically reliable (as the ambient dimension increases) at SNR levels significantly lower than non-adaptive limits. The distilled sensing idea was recently extended to a more general setting of sequential multiple hypothesis testing in [11]; in [12] it is shown that a sequential thresholding procedure comes within a small factor of the optimal sequential procedure in terms of the number of observations needed for asymptotically exact support recovery.

In the second class of models, the observations can consist of arbitrary linear combinations, as in compressed sensing, but for the most part the resource budget is assumed to be discrete, measured in units of normalized observations ([11], [12] also assume a discrete budget). In [13], the distilled sensing approach was extended to the compressed measurement setting. In [14], [15], a Bayesian signal model is adopted and each new observation is chosen to approximately maximize the information gain; [15] is computationally simpler but is most suited to signals with a single non-zero component, i.e., 1-sparse signals. Others have also taken the approach of decomposing the problem into subproblems involving 1-sparse signals and then applying a form of bisection search [16], [17], [18]; [18] employs a more sophisticated search in which the rate of division accelerates, reducing the complexity to doubly logarithmic instead of merely logarithmic. The adaptive methods in [17], [18] were shown to require fewer measurements than the best non-adaptive method. In [16] and [18] however, noise is either not

considered or not fully taken into account. Somewhat different from the aforementioned works is [19], which describes a compressed sensing method that is sequential in the sense that it terminates once the reconstruction error is determined to have fallen below a threshold, but the form of the measurements is not adapted during the process.

Adaptive sensing and resource allocation have also been applied to other classes of signals with more structure. Tree-structured sparsity is considered in [20], which proposes selective sampling of wavelet coefficients based on already sampled coefficients nearby and at coarser resolutions. For two-dimensional piecewise-constant signals, a method that concentrates measurements near boundaries is presented and analyzed in [21], [22]. Adaptive waveform amplitude design is investigated in [23] for unstructured (i.e. dense) parameter estimation in a linear Gaussian model under an average energy constraint.

This paper addresses the problem of estimation and adaptive resource allocation under the first observation model in which components are measured directly. We extend the two-stage allocation policy in [3] to an arbitrary number of stages, focusing on estimation error explicitly as contrasted with performance bounds in [3]. Our method is computationally tractable for a wide range of estimation loss functions satisfying a mild convexity condition, including such commonly used criteria as mean squared error (MSE) and mean absolute error (MAE). The observation model in [3], [10] is also generalized by allowing the observation precision to depend on an arbitrary concave function of the sensing effort. It is shown that the problem can be formulated as a dynamic program, a framework that facilitates the development of allocation policies. An approximate dynamic programming solution is proposed based on open-loop feedback control (OLFC). The performance of these OLFC policies improves monotonically with the number of stages, and in particular improves upon optimal two-stage policies including the one in [3]. Numerical simulations show error reductions up to 4.5 dB relative to the optimal two-stage policy and dramatic reductions relative to non-adaptive sensing, approaching the oracle limit at high SNR. The OLFC policies are also shown to outperform distilled sensing [10] at all SNR and most significantly at higher SNR. The advantages carry over to a radar imaging example that challenges some of the assumptions of our model.

The remainder of the paper proceeds as follows. In Section II, the signal and observation models are specified and a problem of resource-constrained sequential estimation is formulated and then recast as a dynamic program. In Section III, optimal and OLFC approaches to the problem are discussed and a family of OLFC policies is proposed. Numerical simulations comparing our OLFC policies to other policies are presented in Section IV. In Section V, an application to radar imaging is described. Conclusions and future directions are given in Section VI.

II. PROBLEM FORMULATION

We consider signals that are sparse with respect to the standard basis (without loss of generality) in an N -dimensional space. The signal support is represented by a set of indicators I_i , $i = 1, \dots, N$, with $I_i = 1$ if i is in the support and $I_i = 0$ otherwise. We use a probabilistic model in which $I_i = 1$ with prior probability $p_i(0)$, independently of the other indicators. The non-zero signal amplitudes are modelled as independent Gaussian random variables θ_i with prior means $\mu_i(0)$ and variances $\sigma_i^2(0)$. As in [3], [4], a non-informative uniform prior is assumed with $p_i(0) = p_0$, $\mu_i(0) = \mu_0$, and $\sigma_i^2(0) = \sigma_0^2$ for all i , although the theory developed below could also accommodate non-uniform priors.

A sequence of T observations are made with non-negative effort levels $\lambda_i(t)$ that can vary with index i and time $t = 0, \dots, T - 1$. Depending on the application, the effort $\lambda_i(t)$ might represent observation time, number of samples, energy, cost, or computation. It is assumed that the precision (inverse variance) of an observation varies with effort according to a non-decreasing function h . Specifically, given $\lambda_i(t-1)$, the observation of the i th component at time t takes the form

$$y_i(t) = I_i \theta_i + \frac{n_i(t)}{\sqrt{h(\lambda_i(t-1))}}, \quad i = 1, \dots, N, \quad t = 1, \dots, T, \quad (1)$$

where $n_i(t)$ represents i.i.d. zero-mean Gaussian noise with variance σ^2 . The function h is often linear, but nonlinear dependences can also arise. For example, the sensing system may contain nonlinear components such as amplifiers, or the observations may result from integrating a continuous-time random process over an interval of length $\lambda_i(t-1)$ and the process exhibits short-term correlation. In any case, it is reasonable to assume that $h(0) = 0$ and $h(\lambda) > 0$ for $\lambda > 0$, and to normalize h so that $h(1) = 1$. We restrict our attention to static signals so that the signal component $I_i \theta_i$ in (1) does not change with time. For convenience, we use the vector notation $\mathbf{y}(t) = [y_1(t) \dots y_N(t)]^T$ (similarly for other indexed quantities) and denote by $\mathbf{Y}(t) = \{\mathbf{y}(1), \dots, \mathbf{y}(t)\}$ the history of observations up to time t .

The task is to determine the distribution of sensing effort over components and time subject to a total budget constraint,

$$\sum_{t=0}^{T-1} \sum_{i=1}^N \lambda_i(t) = \Lambda_0. \quad (2)$$

Under the normalization $\Lambda_0 = N$, each component receives an average of one unit of effort over time. In the case of single-stage non-adaptive sensing ($T = 1$) and a uniform prior, the most natural choice is to set $\lambda_i(0) = 1$ for all i . Thus σ^2 can be regarded as the noise variance realized under a non-adaptive uniform allocation. In multistage adaptive sensing, the allocation $\boldsymbol{\lambda}(t)$ at time t can depend on the observations $\mathbf{Y}(t)$ collected up to that point. This information allows more resources to be focused on the region of

signal support, thereby improving the SNR. The mapping from $\mathbf{Y}(t)$ to $\lambda(t)$ is referred to as an effort allocation policy. For notational brevity, we will not make the dependence of $\lambda(t)$ on $\mathbf{Y}(t)$ explicit.

To select the effort allocation policy, we seek to minimize the sum over the signal support of the expected losses associated with estimates $\hat{\theta}_i$ of the amplitudes θ_i , based on all observations up to time T :

$$\mathbf{E} \left\{ \sum_{i=1}^N I_i L \left(\left| \hat{\theta}_i - \theta_i \right| \right) \right\}, \quad (3)$$

where the expectation is taken over \mathbf{I} , $\boldsymbol{\theta}$, and $\mathbf{Y}(T)$. The loss function L is assumed to be non-decreasing (and symmetric given the form in (3)). To relate the expected cost (3) to the effort allocation policy, we nest the expectations in the order $\mathbf{Y}(T)$, \mathbf{I} , $\boldsymbol{\theta}$ (outer to inner) and expand to yield

$$\mathbf{E}_{\mathbf{Y}(T)} \left\{ \sum_{i=1}^N p_i(T) \mathbf{E} \left[L \left(\left| \hat{\theta}_i - \theta_i \right| \right) \mid I_i = 1, \mathbf{Y}(T) \right] \right\}, \quad (4)$$

where we have defined $p_i(t) = \Pr(I_i = 1 \mid \mathbf{Y}(t))$. It is shown in Appendix A that $\theta_i \mid I_i = 1, \mathbf{Y}(t)$ has a Gaussian distribution for all t with mean $\mu_i(t)$ and variance $\sigma_i^2(t)$. Since the Gaussian distribution is symmetric and strictly unimodal and the loss function L is also symmetric and non-decreasing, the inner expectation in (4) is minimized by choosing $\hat{\theta}_i = \mu_i(T)$ (a proof of this fact is given in Appendix B). Then the minimum value of the inner expectation depends only on $\sigma_i^2(T)$ and (4) can be expressed as

$$2 \mathbf{E}_{\mathbf{Y}(T)} \left\{ \sum_{i=1}^N p_i(T) \int_0^\infty L(\sigma_i(T)\theta) \phi(\theta) d\theta \right\}, \quad (5)$$

where $\phi(\theta)$ denotes the standard Gaussian probability density function (pdf). The final-stage variance $\sigma_i^2(T)$ depends in turn on the effort allocation according to the relation

$$\sigma_i^2(T) = \frac{\sigma^2}{\sigma^2/\sigma_0^2 + \sum_{t=0}^{T-1} h(\lambda_i(t))}, \quad (6)$$

which is derived in Appendix A. In summary, the problem is to minimize the expected cost defined by (5) and (6) with respect to the effort allocation policy $\lambda(0), \dots, \lambda(T-1)$, subject to the total effort constraint (2).

In the case of the square loss $L(a) = a^2$, i.e., the mean squared error (MSE) criterion, the integral in (5) can be evaluated to yield $\sigma_i^2(T)$, thus reducing (5) to

$$\sigma^2 \mathbf{E}_{\mathbf{Y}(T)} \left\{ \sum_{i=1}^N \frac{p_i(T)}{\sigma^2/\sigma_0^2 + \sum_{t=0}^{T-1} h(\lambda_i(t))} \right\}. \quad (7)$$

The cost function in (7) is closely related to the cost function in [3] although the motivations differ with the latter being related to Chernoff and Cramér-Rao bounds on detection and estimation performance respectively. The general form of the cost function in [3] can be obtained from (7) by replacing $p_i(T)$

with the weighted average $\nu p_i(T) + (1-\nu)(1-p_i(T))$ for $\nu \in [1/2, 1]$, letting $\sigma_0^2 \rightarrow \infty$ so that $\sigma^2/\sigma_0^2 \rightarrow 0$, and choosing h to be the identity function. Given that the generalization of $p_i(T)$ to a weighted average is straightforward to accommodate, we keep $\nu = 1$ to simplify notation in the remainder of the paper.

A. Formulation as a dynamic program

The determination of an optimal effort allocation policy according to (5) and (6) can be formulated as a dynamic program. Although the dynamic programming viewpoint does not offer significant simplifications, it does make available a well-developed set of approaches to the problem, some of which are considered in Section III. Further background in dynamic programming can be found in [24].

To formulate a sequential decision problem as a dynamic program, the cost function must be expressible as a sum of terms indexed by time t , where each term depends only on the current system state $\mathbf{x}(t)$ and the current control action, in our case the effort allocation $\lambda(t)$ (each term may also depend on a random disturbance but this is not required here). The cost function (5) can be recast in the required time-separable form by defining the state $\mathbf{x}(t)$ as $\mathbf{x}(t) = (\mathbf{p}(t), \boldsymbol{\mu}(t), \boldsymbol{\sigma}^2(t), \Lambda(t))$, where $\Lambda(t)$ represents the effort budget remaining at time t . The state variables are initialized as $p_i(0) = p_0$, $\mu_i(0) = \mu_0$, $\sigma_i^2(0) = \sigma_0^2$, and $\Lambda(0) = \Lambda_0$, and evolve according to the following recursions derived in Appendix A:

$$p_i(t+1) = \frac{p_i(t)\phi_1}{p_i(t)\phi_1 + (1-p_i(t))\phi_0}, \quad (8a)$$

$$\mu_i(t+1) = \frac{\sigma^2 \mu_i(t) + h(\lambda_i(t))\sigma_i^2(t)y_i(t+1)}{\sigma^2 + h(\lambda_i(t))\sigma_i^2(t)}, \quad (8b)$$

$$\sigma_i^2(t+1) = \frac{\sigma^2 \sigma_i^2(t)}{\sigma^2 + h(\lambda_i(t))\sigma_i^2(t)}, \quad (8c)$$

$$\Lambda(t+1) = \Lambda(t) - \sum_{i=1}^N \lambda_i(t), \quad (8d)$$

where

$$\phi_0 = \phi(y_i(t+1); 0, \sigma^2/h(\lambda_i(t))),$$

$$\phi_1 = \phi(y_i(t+1); \mu_i(t), \sigma_i^2(t) + \sigma^2/h(\lambda_i(t))),$$

and $\phi(\cdot; \mu, \sigma^2)$ denotes a Gaussian pdf with mean μ and variance σ^2 .

Given the above state definition, we use (8c) to rewrite the denominator in (6) as

$$\frac{\sigma^2}{\sigma_0^2} + \sum_{t=0}^{T-1} h(\lambda_i(t)) = \frac{\sigma^2}{\sigma_i^2(T-1)} + h(\lambda_i(T-1)). \quad (9)$$

We then decompose $\mathbf{E}_{\mathbf{Y}(T)}$ as $\mathbf{E}_{\mathbf{Y}(T-1)} \mathbf{E}_{\mathbf{Y}(T)|\mathbf{Y}(T-1)}$ and notice that only $p_i(T)$ in (5) depends on $\mathbf{y}(T)$. Taking the expectation of (34) with respect to $\mathbf{y}(t) | \mathbf{Y}(t-1)$, we deduce that

$$\mathbf{E}_{\mathbf{y}(t)}\{p_i(t) | \mathbf{Y}(t-1)\} = p_i(t-1), \quad t = 1, \dots, T. \quad (10)$$

Using (6), (9), and (10), the effort allocation problem may be stated as

$$\begin{aligned} \min_{\lambda(0), \dots, \lambda(T-1)} \quad & \mathbf{E}_{\mathbf{Y}(T-1)} \{G(\mathbf{x}(T-1), \boldsymbol{\lambda}(T-1))\} \\ \text{s.t.} \quad & \sum_{t=0}^{T-1} \sum_{i=1}^N \lambda_i(t) = \Lambda_0, \quad \lambda_i(t) \geq 0 \quad \forall t, i, \end{aligned} \quad (11)$$

where the cost function is of the desired form with a single non-zero term at time $T-1$,

$$G(\mathbf{x}(T-1), \boldsymbol{\lambda}(T-1)) = \sum_{i=1}^N p_i(T-1) g(\sigma_i^2(T-1), h(\lambda_i(T-1))), \quad (12)$$

$$g(\sigma_i^2(t), \bar{h}_i) = \int_0^\infty L \left(\frac{\sigma\theta}{\sqrt{\sigma^2/\sigma_i^2(t) + \bar{h}_i}} \right) \phi(\theta) d\theta, \quad (13)$$

depending explicitly on $\mathbf{p}(T-1)$, $\boldsymbol{\sigma}^2(T-1)$, and $\boldsymbol{\lambda}(T-1)$. The dependence on the variables $\boldsymbol{\lambda}(t)$, $t = 0, \dots, T-2$ is implicit through the probability distribution of the observations $\mathbf{Y}(T-1)$ and the recursions in (8). The constraints in (11) actually represent a continuum of constraints since they are required to be satisfied for all realizations of $\mathbf{Y}(T-1)$.

III. EFFORT ALLOCATION POLICIES

In this section, we develop policies directed at solving the effort allocation problem (11). Optimal policies are discussed in Section III-A while a less complex method known as open-loop feedback control is discussed in Section III-B. We then discuss two approaches to improving the performance of OLFC: generalized OLFC in Section III-C, and policy rollout in Section III-D.

A. Optimal policies

In principle, it is possible to employ exact dynamic programming to determine an optimal policy for (11). The dynamic programming approach decomposes (11) into a sequence of optimizations proceeding backward in time, making use of the chain rule $\mathbf{E}_{\mathbf{Y}(T-1)} = \mathbf{E}_{\mathbf{Y}(1)} \mathbf{E}_{\mathbf{Y}(2)|\mathbf{Y}(1)} \dots \mathbf{E}_{\mathbf{Y}(T-1)|\mathbf{Y}(T-2)}$ and the fact that each allocation $\boldsymbol{\lambda}(t)$ is a function of past observations $\mathbf{Y}(t)$ but not future ones. The last-stage

optimization is given by

$$\begin{aligned}
J_{T-1}^*(\mathbf{x}(T-1)) &= \min_{\boldsymbol{\lambda}(T-1)} G(\mathbf{x}(T-1), \boldsymbol{\lambda}(T-1)) \\
\text{s.t.} \quad &\sum_{i=1}^N \lambda_i(T-1) = \Lambda(T-1), \\
&\lambda_i(T-1) \geq 0 \quad \forall i,
\end{aligned} \tag{14a}$$

and for $t = T-2, T-3, \dots, 0$, the optimizations are defined recursively as follows:

$$\begin{aligned}
J_t^*(\mathbf{x}(t)) &= \min_{\boldsymbol{\lambda}(t)} \mathbf{E}_{\mathbf{y}(t+1)} \{J_{t+1}^*(\mathbf{x}(t+1)) \mid \mathbf{x}(t), \boldsymbol{\lambda}(t)\} \\
\text{s.t.} \quad &\sum_{i=1}^N \lambda_i(t) \leq \Lambda(t), \quad \lambda_i(t) \geq 0 \quad \forall i.
\end{aligned} \tag{14b}$$

The desired optimal cost in (11) is $J_0^*(\mathbf{x}(0))$. The notation in (14b) reflects the fact that the distribution of $\mathbf{y}(t+1)$ given $\mathbf{Y}(t)$ is completely determined by $\mathbf{x}(t)$ and $\boldsymbol{\lambda}(t)$; more specifically, $f(y_i(t+1) \mid \mathbf{Y}(t))$ is given by the denominator of the right-hand side of (8a) as can be seen from (36). The next state $\mathbf{x}(t+1)$ is specified by $\mathbf{x}(t)$, $\boldsymbol{\lambda}(t)$, and $\mathbf{y}(t+1)$ through (8). Thus the choice of $\boldsymbol{\lambda}(t)$ depends on $\mathbf{Y}(t)$ only through the state $\mathbf{x}(t)$, which is a property of dynamic programs [24].

An optimal policy can be obtained by first solving (14a) for $\boldsymbol{\lambda}(T-1)$ and then using the result in (14b) to solve for $\boldsymbol{\lambda}(T-2)$. The remaining allocations are determined in the same recursive way. This exact procedure is computationally tractable in a few cases. For $T = 1$, it suffices to solve (14a), which is a convex optimization problem under some conditions to be discussed in Section III-B. For $T = 2$ and a uniform prior ($p_i(0) = p_0$, $\mu_i(0) = \mu_0$, $\sigma_i^2(0) = \sigma_0^2$), symmetry allows the initial allocation $\boldsymbol{\lambda}(0)$ to be restricted to the form $\boldsymbol{\lambda}(0) = \beta^{(2)}(0)\mathbf{1}$, where $\mathbf{1}$ denotes a vector with unit entries. Thus (14b) becomes a one-dimensional optimization with respect to the multiplier $\beta^{(2)}(0)$. For fixed $\beta^{(2)}(0)$, the expectation in (14b) can be evaluated by sampling from the distribution of $\mathbf{y}(1)$ and then solving (14a) for the resulting values of the state $\mathbf{x}(1)$.

For $T > 2$ however, an exact solution via (14a) and (14b) is very difficult. Problem (14b) is in general an N -dimensional optimization with no known structure and N potentially large. Furthermore, a common approach in which the cost-to-go $J_{t+1}^*(\mathbf{x}(t+1))$ in (14b) is summarized by its values at a small number of representative states $\mathbf{x}(t+1)$ is made difficult by the high dimension ($3N+1$ to be exact) and continuous nature of the state. For these reasons, we do not consider an exact solution to (11) for $T > 2$, opting instead for an approximate method as is discussed next.

B. Open-loop feedback control

A well-known approach to approximate dynamic programming is that of open-loop feedback control (OLFC) [24]. We consider the problem of determining the allocation $\lambda(t)$ at time t given the current set of observations $\mathbf{Y}(t)$, or equivalently the state $\mathbf{x}(t)$. In OLFC, this computation is simplified by assuming that future allocations $\lambda(t+1), \dots, \lambda(T-1)$ can depend only on $\mathbf{Y}(t)$ and not future observations. In light of this assumption, the only quantities that depend on $\mathbf{y}(t+1), \dots, \mathbf{y}(T-1)$ in (12) are the probabilities $p_i(T-1)$. The expectations $\mathbf{E}_{\mathbf{y}(t+1)|\mathbf{Y}(t)} \dots \mathbf{E}_{\mathbf{y}(T-1)|\mathbf{Y}(T-2)}$ in (11) can then be applied to transform $p_i(T-1)$ into $p_i(t)$ using (10) repeatedly. The resulting cost function is to be optimized with respect to $\lambda(t), \dots, \lambda(T-1)$ jointly, leading to the problem

$$\begin{aligned} \min_{\lambda(t), \dots, \lambda(T-1)} \quad & \sum_{i=1}^N p_i(t) g \left(\sigma_i^2(t), \sum_{\tau=t}^{T-1} h(\lambda_i(\tau)) \right) \\ \text{s.t.} \quad & \sum_{\tau=t}^{T-1} \sum_{i=1}^N \lambda_i(\tau) = \Lambda(t), \quad \lambda_i(\tau) \geq 0 \quad \forall \tau, i, \end{aligned} \quad (15)$$

where we have made use of a rearrangement similar to (9). The budget constraint in (15) is assumed to be met with equality as otherwise the cost could be decreased.

For $t = T-1$, the OLFC problem (15) coincides with the last-stage optimization in (14a). For $t < T-1$, OLFC represents a significant simplification relative to the exact optimization in (14b) because the cost function in (15) is expressed explicitly without the need to evaluate expectations recursively. Under certain conditions specified in the following proposition, problem (15) is also a convex optimization and thus can be tractably solved.

Proposition 1: The OLFC problem (15) is a convex optimization problem if the loss function L is non-decreasing, $g(\sigma_i^2(t), \bar{h}_i)$ in (13) is a convex function of \bar{h}_i for $\bar{h}_i \geq 0$ and all $\sigma_i^2(t)$, and the effort function h is concave.

Proof: Since the constraints in (15) are all linear, the feasible set is convex (more precisely a simplex). The cost function is a non-negative combination of functions $g(\sigma_i^2(t), \bar{h}_i)$ with $\bar{h}_i = \sum_{\tau=t}^{T-1} h(\lambda_i(\tau))$, so it suffices to prove that g is convex as a function of $\lambda_i(t), \dots, \lambda_i(T-1)$. First note that \bar{h}_i , as a sum of concave functions, is concave in $\lambda_i(t), \dots, \lambda_i(T-1)$. Given that L is a non-decreasing function of its argument, g is seen to be a non-increasing function of \bar{h}_i . Furthermore, we may extend the definition of g to negative \bar{h}_i by letting $g(\sigma_i^2(t), \bar{h}_i) = \infty$ for $\bar{h}_i < 0$, thereby preserving the monotonicity and assumed convexity of g . It then follows from a property of compositions of functions [25] that g is convex in $\lambda_i(t), \dots, \lambda_i(T-1)$. \blacksquare

The assumptions in Proposition 1 are not difficult to satisfy. It was already assumed in Section II that

L is non-decreasing so that the optimal amplitude estimate $\hat{\theta}_i$ is equal to the conditional mean $\mu_i(T)$. The concavity assumption on h is satisfied by the identity function as well as functions corresponding to a sublinear dependence of the observation precision on sensing effort. The convexity assumption on g is satisfied by a variety of commonly used loss functions. As a first example we consider the 0-1 loss function for a tolerance ϵ ,

$$L_\epsilon(a) = \begin{cases} 0, & 0 \leq a < \epsilon, \\ 1, & a \geq \epsilon. \end{cases}$$

The integral in (13) may be evaluated in this case to yield

$$g(\sigma_i^2(t), \bar{h}_i) = Q\left(\frac{\epsilon}{\sigma} \sqrt{\sigma^2/\sigma_i^2(t) + \bar{h}_i}\right),$$

where Q denotes the Q-function, i.e., the standard Gaussian tail probability. Since the Q-function is convex decreasing for non-negative arguments and the square root function is concave in \bar{h}_i , the same property used in the proof of Proposition 1 may be invoked to conclude that g is a convex function of \bar{h}_i . The convexity of g can also be verified for $L(a) = 1 - e^{-ba}$ with $b > 0$, which can be regarded as a continuous approximation to the 0-1 loss function.

The assumption that g is convex may be replaced by one of the following stricter but more easily checked conditions:

- (a) $L\left(\sigma/\sqrt{\sigma^2/\sigma_i^2(t) + \bar{h}_i}\right)$ is a convex function of \bar{h}_i ;
- (b) L is convex.

Condition (a) implies that g is convex because scaling the argument of L by θ does not affect convexity and because the weighting function $\phi(\theta)$ in the integral in (13) is always positive. Condition (b) implies condition (a) because of a composition property similar to the one used earlier and the convexity of $\sigma/\sqrt{\sigma^2/\sigma_i^2(t) + \bar{h}_i}$ with respect to \bar{h}_i . If L is twice differentiable, condition (a) can be shown to be equivalent to the inequality

$$aL''(a) + 3L'(a) \geq 0, \quad a \geq 0, \tag{16}$$

whereas (b) is equivalent to $L''(a) \geq 0$. Condition (b) includes the square loss $L(a) = a^2$ corresponding to MSE, the linear loss $L(a) = a$ corresponding to mean absolute error (MAE), the Huber loss which combines the square and linear losses in a continuous and convex manner, and the two-sided hinge loss. More generally, (16) is satisfied for any power-law function $L(a) = a^q$ with $q > 0$ and for $L(a) = \log(1 + ba)$ with $b > 0$. Note that a^q for $0 < q < 1$ and $\log(1 + ba)$ are concave functions of a . Taking the limit as $q \rightarrow 0$ of the power-law functions yields the 0-1 loss function, which was shown earlier to result in a convex g .

In the remainder of the paper, we assume that the assumptions of Proposition 1 are satisfied and hence the OLFC problem (15) is a convex optimization. We now address the solution of (15). The cost function in (15) depends on $\lambda(t), \dots, \lambda(T-1)$ only through the quantities $\bar{h}_i = \sum_{\tau=t}^{T-1} h(\lambda_i(\tau))$, and is more specifically a non-increasing function of \bar{h}_i as argued in the proof of Proposition 1. Therefore (15) may be solved via a two-step procedure: first we fix $\bar{\lambda}_i(t) = \sum_{\tau=t}^{T-1} \lambda_i(\tau)$ and seek to maximize \bar{h}_i as functions of $\bar{\lambda}_i(t)$, i.e.,

$$\begin{aligned} \bar{h}_i^*(\bar{\lambda}_i(t)) = \max_{\lambda_i(t), \dots, \lambda_i(T-1)} & \sum_{\tau=t}^{T-1} h(\lambda_i(\tau)) \\ \text{s.t.} & \sum_{\tau=t}^{T-1} \lambda_i(\tau) = \bar{\lambda}_i(t), \\ & \lambda_i(\tau) \geq 0 \quad \forall \tau, i, \end{aligned} \quad (17)$$

and then we substitute the maximum values $\bar{h}_i^*(\bar{\lambda}_i(t))$ into (15) and optimize with respect to $\bar{\lambda}_i(t)$. The maximum $\bar{h}_i^*(\bar{\lambda}_i(t))$ can be determined by noting that (17) is a concave maximization problem subject to a simplex constraint. For such problems, we have the following necessary and sufficient optimality condition:

$$\text{if } \lambda_i^*(\tau) > 0 \quad \text{then} \quad \frac{\partial \bar{h}_i}{\partial \lambda_i(\tau)} \geq \frac{\partial \bar{h}_i}{\partial \lambda_i(\tau')} \quad \forall \tau' \neq \tau, \quad (18)$$

where the partial derivatives are evaluated at the optimum. The solution $\lambda_i^*(\tau) = \bar{\lambda}_i(t)/(T-t)$ for all τ satisfies (18) by symmetry since all of the partial derivatives are equal. The corresponding maximum value is therefore $\bar{h}_i^*(\bar{\lambda}_i(t)) = (T-t)h(\bar{\lambda}_i(t)/(T-t))$. Note however that the optimal solution to (17) may not be unique if h is not strictly concave. In particular, if h is the identity function, then $\bar{h}_i = \bar{\lambda}_i(t)$ regardless of the choice of $\lambda_i(t), \dots, \lambda_i(T-1)$. We return to the issue of non-uniqueness in Section III-C.

With the substitutions $\sum_{\tau=t}^{T-1} \lambda_i(\tau) = \bar{\lambda}_i(t)$ and $\sum_{\tau=t}^{T-1} h(\lambda_i(\tau)) = (T-t)h(\bar{\lambda}_i(t)/(T-t))$, (15) simplifies to

$$\begin{aligned} \min_{\bar{\lambda}(t)} & \sum_{i=1}^N p_i(t) g \left(\sigma_i^2(t), (T-t)h \left(\frac{\bar{\lambda}_i(t)}{T-t} \right) \right) \\ \text{s.t.} & \sum_{i=1}^N \bar{\lambda}_i(t) = \Lambda(t), \quad \bar{\lambda}_i(t) \geq 0 \quad \forall i, \end{aligned} \quad (19)$$

a simplex-constrained convex minimization problem. Problem (19) thus satisfies an optimality condition similar to (18) with the inequality between partial derivatives reversed in direction. This condition implies that optimal solutions to (19) have certain properties akin to water-filling. First, the solutions exhibit thresholding in the sense that $\bar{\lambda}_i^*(t)$ must be zero if the corresponding partial derivative is not among

the lowest. Second, the partial derivatives corresponding to non-zero components must all be equal. This in turn induces an ordering among the non-zero allocations as a function of the probabilities $p_i(t)$ and variances $\sigma_i^2(t)$.

To illustrate the properties of optimal solutions to (19), we specialize to the case of power-law losses $L(a) = a^q$ and the identity effort function $h(\lambda) = \lambda$. In this case, (19) reduces to

$$\begin{aligned} \min_{\bar{\lambda}(t)} \quad & \sum_{i=1}^N \frac{p_i(t)}{(\sigma^2/\sigma_i^2(t) + \bar{\lambda}_i(t))^{q/2}} \\ \text{s.t.} \quad & \sum_{i=1}^N \bar{\lambda}_i(t) = \Lambda(t), \quad \bar{\lambda}_i(t) \geq 0 \quad \forall i, \end{aligned} \quad (20)$$

and the optimal solution can be stated explicitly. A detailed derivation is provided in Appendix C. First we define $\gamma = 2/(q + 2)$ and π to be an index permutation that sorts the quantities $p_i^\gamma(t)\sigma_i^2(t)$ in non-increasing order:

$$p_{\pi(1)}^\gamma(t)\sigma_{\pi(1)}^2(t) \geq p_{\pi(2)}^\gamma(t)\sigma_{\pi(2)}^2(t) \geq \dots \geq p_{\pi(N)}^\gamma(t)\sigma_{\pi(N)}^2(t). \quad (21)$$

Next define $b(k)$ to be the monotonically non-decreasing function of $k = 0, 1, \dots, N$ with $b(N) = \infty$ and

$$b(k) = \frac{\sigma^2}{p_{\pi(k+1)}^\gamma(t)\sigma_{\pi(k+1)}^2(t)} \sum_{i=1}^k p_{\pi(i)}^\gamma(t) - \sum_{i=1}^k \frac{\sigma^2}{\sigma_{\pi(i)}^2(t)}, \quad k = 0, \dots, N-1. \quad (22)$$

Then the optimal solution $\bar{\lambda}^*(t)$ to (20) is given by

$$\bar{\lambda}_{\pi(i)}^*(t) = \begin{cases} Cp_{\pi(i)}^\gamma(t) - \frac{\sigma^2}{\sigma_{\pi(i)}^2(t)}, & i = 1, \dots, k, \\ 0, & i = k+1, \dots, N, \end{cases} \quad (23)$$

where

$$C = \frac{\Lambda(t) + \sum_{j=1}^k \frac{\sigma^2}{\sigma_{\pi(j)}^2(t)}}{\sum_{j=1}^k p_{\pi(j)}^\gamma(t)} \quad (24)$$

and the number of non-zero components k is determined by the interval $(b(k-1), b(k)]$ to which the budget parameter $\Lambda(t)$ belongs. The monotonicity of $b(k)$ ensures that the mapping from $\Lambda(t)$ to k is well-defined. We note that k and C could also be computed using the general procedure in [26]. The thresholding property is clearly seen in (23). Furthermore, the non-zero allocations increase with the probabilities $p_i(t)$ raised to the power γ and decrease with the precisions $1/\sigma_i^2(t)$.

In the case of general loss and effort functions, (19) may not have an explicit solution as in (21)–(24). Nevertheless, an efficient iterative solution is possible under the assumption of convexity. One possibility is to use a projected gradient algorithm, taking advantage of the ease of projecting onto a simplex.

The solution to (19) specifies the values of the sums $\bar{\lambda}_i(t) = \sum_{\tau=t}^{T-1} \lambda_i(\tau)$. However, the solution to (17) may not uniquely specify the division of $\bar{\lambda}_i(t)$ into $\lambda_i(t), \dots, \lambda_i(T-1)$ if the effort function h is not strictly concave. In addition, since the OLFC optimization (19) is similar to the last-stage optimization (14a), the resulting policy can be somewhat aggressive in allocating effort to components currently believed to contain signal as opposed to waiting for further confirmation. In the next subsection, these issues are addressed through a generalization of the OLFC approach.

C. Generalized open-loop feedback control

In this subsection, we discuss two modifications to the OLFC policy in Section III-B. The first modification is directed at optimizing the distribution of effort over stages and applies to all loss functions. The second modification reduces premature exploitation and is presented only for power-law loss functions; similar strategies could be devised for other loss functions.

To optimize the allocation of effort over stages, we restrict the allocation for the current stage $\lambda(t)$ to be proportional to the optimal solution $\bar{\lambda}^*(t)$ of (19), i.e., $\lambda(t) = \beta^{(T)}(t)\bar{\lambda}^*(t)$, where $\beta^{(T)}(t) \in [0, 1]$ represents the fraction of the remaining budget $\Lambda(t)$ that is used at time t and the superscript T denotes the total number of stages. The fractions $\beta^{(T)}(t)$ are chosen based on a generalization of the optimal policies for $T = 1$ and $T = 2$ in Section III-A. Both of these optimal policies belong to the OLFC class. Specifically, the $T = 1$ policy results from solving (14a), which is a special case of (19) with $t = T - 1$, and setting $\beta^{(1)}(0) = 1$ since there is only one stage. The $T = 2$ policy uses an initial allocation $\lambda(0) = \beta^{(2)}(0)\mathbf{1}$, which is of the same form as the solution to (19) for $t = 0$ under a uniform prior, followed by the solution to (19) for $t = 1$ scaled by $\beta^{(2)}(1) = 1$. Note that the second stage in the $T = 2$ policy is identical to the $T = 1$ policy with $\beta^{(2)}(1) = \beta^{(1)}(0)$. For $T > 2$, we follow the same strategy of reusing the $(T - 1)$ -stage fractions in the T -stage policy, setting $\beta^{(T)}(t) = \beta^{(T-1)}(t - 1)$ for $t = 1, 2, \dots, T - 1$. The first-stage fraction $\beta^{(T)}(0)$ is then optimized as described below.

The second modification is to allow the exponent γ in (21)–(24) to vary with time. The last-stage exponent $\gamma^{(T)}(T - 1)$ is set to $2/(q + 2)$, the optimal exponent for the loss function $L(a) = a^q$. In earlier stages, smaller exponents are used to make the policy more conservative, specifically by weakening the dependence on the probabilities $p_i(t)$. We propose the simple strategy of optimizing only the first-stage exponent $\gamma^{(T)}(0)$ and constraining the remaining exponents to linearly interpolate between $\gamma^{(T)}(0)$ and $\gamma^{(T)}(T - 1)$. This reduces the determination of the fractions $\beta^{(T)}(t)$ and exponents $\gamma^{(T)}(t)$ to a two-dimensional optimization regardless of the number of stages.

The first-stage parameters $\beta^{(T)}(0)$ and $\gamma^{(T)}(0)$ are determined recursively for $T = 1, 2, \dots$ starting

from $\beta^{(1)}(0) = 1$ and $\gamma^{(1)}(0) = 2/(q+2)$. Define $J_t^{(T)}(\mathbf{x}(t))$ to be the cost-to-go of a T -stage policy in this family starting from time t and state $\mathbf{x}(t)$. Then for $T > 1$, $\beta^{(T)}(0)$ and $\gamma^{(T)}(0)$ are given by

$$\left(\beta^{(T)}(0), \gamma^{(T)}(0)\right) = \arg \min_{\substack{0 \leq \beta \leq 1 \\ \gamma \leq 2/(q+2)}} \mathbf{E}_{\mathbf{y}(1)} \left\{ J_1^{(T)}(\mathbf{x}(1)) \mid \mathbf{x}(0), \beta \bar{\boldsymbol{\lambda}}^*(0) \right\}. \quad (25)$$

The parameters $\beta^{(T)}(1), \dots, \beta^{(T)}(T-2)$ and $\gamma^{(T)}(1), \dots, \gamma^{(T)}(T-2)$ required to evaluate $J_1^{(T)}$ are specified by the $(T-1)$ -stage policy and the choice of γ . The expectation in (25) can be computed by sampling from the distribution of $\mathbf{y}(1)$, determining the state $\mathbf{x}(1)$ using (8), and then simulating the remainder of the policy. All of these computations can be done offline since they depend only on the initial state $\mathbf{x}(0)$ and previously determined policies. In addition, since the optimization in (25) can partially account for the effect of future observations on future allocations, an effect that is ignored in the OLFC simplification, the optimization over β is performed even in the case of strictly concave h . Otherwise, (17) would yield a uniform distribution over stages corresponding to $\beta^{(T)}(t) = 1/(T-t)$.

The family of generalized OLFC policies defined above satisfies the following monotonic improvement property.

Proposition 2: The cost of the generalized OLFC policies is non-increasing in the number of stages, i.e.,

$$J_0^{(T)}(\mathbf{x}(0)) \leq J_0^{(T-1)}(\mathbf{x}(0)), \quad T = 2, 3, \dots$$

Proof: The cost of the T -stage policy is given by

$$J_0^{(T)}(\mathbf{x}(0)) = \min_{\substack{0 \leq \beta \leq 1 \\ \gamma \leq 2/(q+2)}} \mathbf{E}_{\mathbf{y}(1)} \left\{ J_1^{(T)}(\mathbf{x}(1)) \mid \mathbf{x}(0), \beta \bar{\boldsymbol{\lambda}}^*(0) \right\}. \quad (26)$$

Consider fixing $\beta = 0$ and

$$\gamma = \begin{cases} \frac{2}{q+2}, & T = 2, \\ \frac{T-1}{T-2} \gamma^{(T-1)}(0) - \frac{1}{T-2} \frac{2}{q+2}, & T > 2 \end{cases}$$

on the right-hand side of (26). With $\beta = 0$, the observations $\mathbf{y}(1)$ are not taken, the state $\mathbf{x}(1)$ is unchanged from $\mathbf{x}(0)$, and the budget usage fractions are the same as in the $(T-1)$ -stage policy. It can also be seen from the choice of γ that the exponents are the same as for $T-1$, and hence the right-hand side of (26) reduces to $J_0^{(T-1)}(\mathbf{x}(0))$. The claim then follows. \blacksquare

Proposition 2 implies in particular that the generalized OLFC policies for $T > 2$ improve upon the optimal policy for $T = 2$. The corresponding performance gains are quantified through numerical simulations in Section IV.

D. Rollout OLFC policies

We now discuss a different approach to improving the performance of OLFC based on the dynamic programming technique of *policy rollout* [24]. For simplicity, we assume that the exponent γ in (21)–(24) is fixed to $2/(q+2)$ in all stages, unlike in Section III-C. In this subsection only, we also make the same assumption for the generalized OLFC policies, i.e., the only parameter optimized in (25) is $\beta^{(T)}(0)$. Rollout could also be applied in the case of time-varying γ by changing the optimization over $\beta(t)$ in (27) below to a joint optimization over $\beta(t)$ and $\gamma(t)$.

In the last stage $t = T - 1$ of a rollout policy, the allocation is determined as before by solving (14a), or equivalently by solving (15) with budget usage fraction $\tilde{\beta}^{(T)}(T - 1) = 1$ (we use a tilde to distinguish the rollout fractions from those in the generalized OLFC policies). For $t = 0, 1, \dots, T - 2$, the fraction $\tilde{\beta}^{(T)}(t)$ is determined according to

$$\tilde{\beta}^{(T)}(t) = \arg \min_{0 \leq \beta(t) \leq 1} \mathbf{E}_{\mathbf{y}(t+1)} \left\{ J_{t+1}^{(T)}(\mathbf{x}(t+1)) \mid \mathbf{x}(t), \beta(t) \bar{\boldsymbol{\lambda}}^*(t) \right\}, \quad (27)$$

where $J_{t+1}^{(T)}(\mathbf{x}(t+1))$ is the cost-to-go of the T -stage generalized policy. Thus $\tilde{\beta}^{(T)}(t)$ is chosen assuming that future stages follow the generalized policy. The corresponding cost-to-go $J_{t+1}^{(T)}(\mathbf{x}(t+1))$ can be viewed as an approximation to the optimal cost-to-go $J_{t+1}^*(\mathbf{x}(t+1))$ in (14b). Comparing (27) with (25) (and assuming that $\gamma(t) = 2/(q+2)$ for all t), it is seen that $\tilde{\beta}^{(T)}(0) = \beta^{(T)}(0)$. In other stages however, the rollout fractions differ from those in the corresponding generalized policy because they are re-optimized based on the value of the current state $\mathbf{x}(t)$ instead of being taken directly from a policy with fewer stages.

In general, rollout policies have the property of improved performance over the policies on which they are based. The same holds for the present rollout policy, with the difference being that the optimization in (27) is restricted to a line search over $\beta(t)$. Denoting by $\tilde{J}_t^{(T)}(\mathbf{x}(t))$ the cost-to-go of a T -stage rollout policy starting from time t and state $\mathbf{x}(t)$, we have the following result:

Proposition 3: The T -stage rollout OLFC policy has a lower cost-to-go than the corresponding generalized OLFC policy in all stages and states, i.e.,

$$\tilde{J}_t^{(T)}(\mathbf{x}(t)) \leq J_t^{(T)}(\mathbf{x}(t)), \quad t = 0, \dots, T - 1, \quad \forall \mathbf{x}(t).$$

Proof: The proof is based on [24, Sec. 6.4]. For $t = T - 1$, the two policies coincide so the costs-to-go are the same. Assume inductively that $\tilde{J}_{t+1}^{(T)}(\mathbf{x}(t+1)) \leq J_{t+1}^{(T)}(\mathbf{x}(t+1))$ for all $\mathbf{x}(t+1)$. The cost-to-go of the rollout policy is defined by

$$\tilde{J}_t^{(T)}(\mathbf{x}(t)) = \mathbf{E}_{\mathbf{y}(t+1)} \left\{ \tilde{J}_{t+1}^{(T)}(\mathbf{x}(t+1)) \mid \mathbf{x}(t), \tilde{\beta}^{(T)}(t) \bar{\boldsymbol{\lambda}}^*(t) \right\}$$

and similarly for the nested policy. By the induction hypothesis and the definition of the rollout policy (27),

$$\begin{aligned}\tilde{J}_t^{(T)}(\mathbf{x}(t)) &\leq \mathbf{E}_{\mathbf{y}(t+1)} \left\{ J_{t+1}^{(T)}(\mathbf{x}(t+1)) \mid \mathbf{x}(t), \tilde{\beta}^{(T)}(t) \bar{\lambda}^*(t) \right\} \\ &\leq \mathbf{E}_{\mathbf{y}(t+1)} \left\{ J_{t+1}^{(T)}(\mathbf{x}(t+1)) \mid \mathbf{x}(t), \beta^{(T)}(t) \bar{\lambda}^*(t) \right\} \\ &= J_t^{(T)}(\mathbf{x}(t))\end{aligned}$$

for all $\mathbf{x}(t)$ as required. Note that the second inequality depends on the generalized policy being included in the class over which the rollout policy is optimized. ■

The rollout OLFC policies can make greater use of knowledge of the state $\mathbf{x}(t)$ but are consequently more demanding computationally than the generalized OLFC policies. Instead of a single optimization in (25), $T - 1$ optimizations as in (27) are required. Furthermore and in contrast to (25), (27) must be evaluated online since it depends on the current state $\mathbf{x}(t)$. The simulations involved in computing the expectation in (27) do become shorter however as t increases. The improvement due to rollout is characterized through numerical simulations in Section IV.

IV. NUMERICAL SIMULATIONS

Numerical simulations are used to evaluate the OLFC policies developed in Section III. The monotonic improvement property of Proposition 2 is verified and gains up to several dB are observed relative to the optimal two-stage policy. The proposed policies are also seen to consistently outperform distilled sensing (DS), most significantly at higher SNR.

In the simulations, we set $N = 10000$ and generate signals and observations according to the model in Section II. Except where indicated, the signal mean μ_0 is normalized to 1 and the signal standard deviation σ_0 is set to $1/4$.

Two families of generalized OLFC policies are considered, one optimized for MSE (final exponent $\gamma^{(T)}(T - 1) = 1/2$, denoted OLFC-MSE) and the other for MAE ($\gamma^{(T)}(T - 1) = 2/3$, denoted OLFC-MAE). The number of stages T is varied from 2 to 10 and the final estimate is given by $\boldsymbol{\mu}^{(T)}$. In the offline determination of the parameters $\beta^{(T)}(t)$ and $\gamma^{(T)}(t)$, the optimization in (25) may be inexact because of finite-sample approximations to the expectations. To mitigate such errors, we make use of the empirical observation that $\beta^{(T)}(0)$ and $\gamma^{(T)}(0)$ appear to vary smoothly with SNR, and $\beta^{(T)}(0)$ also appears to decrease monotonically with T . Accordingly, we first obtain raw estimates of $\beta^{(T)}(0)$ and $\gamma^{(T)}(0)$ and then perform a polynomial fit as a function of SNR, where the polynomials for $\beta^{(T)}(0)$ are

constrained to satisfy $\beta^{(T)}(0) \geq \beta^{(T+1)}(0)$ for all T . In our experience, a polynomial degree of 6 is sufficient to capture the variation of the parameters over the SNR range considered.

For the rollout OLFC policies, the fractions $\tilde{\beta}^{(T)}(t)$ in (27) are also determined through finite-sample approximations to expectations and are thus subject to the same type of error. The difference as noted in Section III-D is that (27) must be evaluated online, and hence the number of samples is limited by computational constraints. To circumvent this tradeoff, we again make use of an empirical smoothness property, this time of the expectation in (27) as a function of $\beta(t)$. Approximations to the expectations are first obtained using a relatively small number of samples, and a fourth-order polynomial in $\beta(t)$ is fit to the approximation. The polynomial fit is then minimized to determine $\tilde{\beta}^{(T)}(t)$. Note that $\beta(t) = 1$ corresponds to a single-stage policy whose cost can be computed exactly from the current state $\mathbf{x}(t)$ as described in Appendix C. Thus $\beta(t) = 1$ and its corresponding single-stage cost represent a fixed point that constrains the polynomial fit.

For DS, while [10] prescribes a single value for T as a function of the dimension N , in our simulations we consider all values of T between 2 and 10 as with OLFC. Following [10], we use a geometrically decreasing allocation of effort over stages with decay ratio $3/4$ and equal first and last stages. More precisely, defining $\alpha(t)$ as the fraction of the total budget used in stage t , we have $\alpha(t) = \alpha(0)(3/4)^t$ for $t = 1, \dots, T-2$, $\alpha(T-1) = \alpha(0)$, and $\alpha(0)$ chosen such that $\sum_{t=0}^{T-1} \alpha(t) = 1$.

In Fig. 1, we plot the MSE and MAE for various policies as a function of SNR, where SNR is defined as $10 \log_{10}(\mu_0^2/\sigma^2)$ in dB. Each point represents the average of 4000 simulations. The baseline corresponding to 0 dB on the vertical axis is the optimal non-adaptive policy, which under a uniform prior allocates one unit of effort to all components. For context, we also include the oracle policy, which distributes effort uniformly over the true signal support. The oracle thus provides an upper bound on the achievable performance, although the bound is unlikely to be tight at lower SNR.

In general, adaptivity yields higher gains for sparser signals ($p_0 = 0.01$) since resources can be concentrated on fewer components once the support is identified. The 10-stage generalized OLFC policies improve upon the 2-stage OLFC policies as expected. The largest gains occur at intermediate SNR and reach 1.5 dB for $p_0 = 0.1$ and 4.5 dB for $p_0 = 0.01$. Recall that the 2-stage OLFC-MSE policy is optimal in terms of MSE for $T = 2$, and similarly for OLFC-MAE. Note also that the performance is only slightly affected by a mismatch between the OLFC policy and the loss function. At high SNR, the OLFC policies approach the oracle gain, which in turn approaches the sparsity factor $1/p_0$. In contrast, the DS policies saturate at significantly lower levels since they are not designed with estimation performance in mind. While the 10-stage DS policy outperforms the optimal 2-stage policy at lower SNR, the 10-stage OLFC

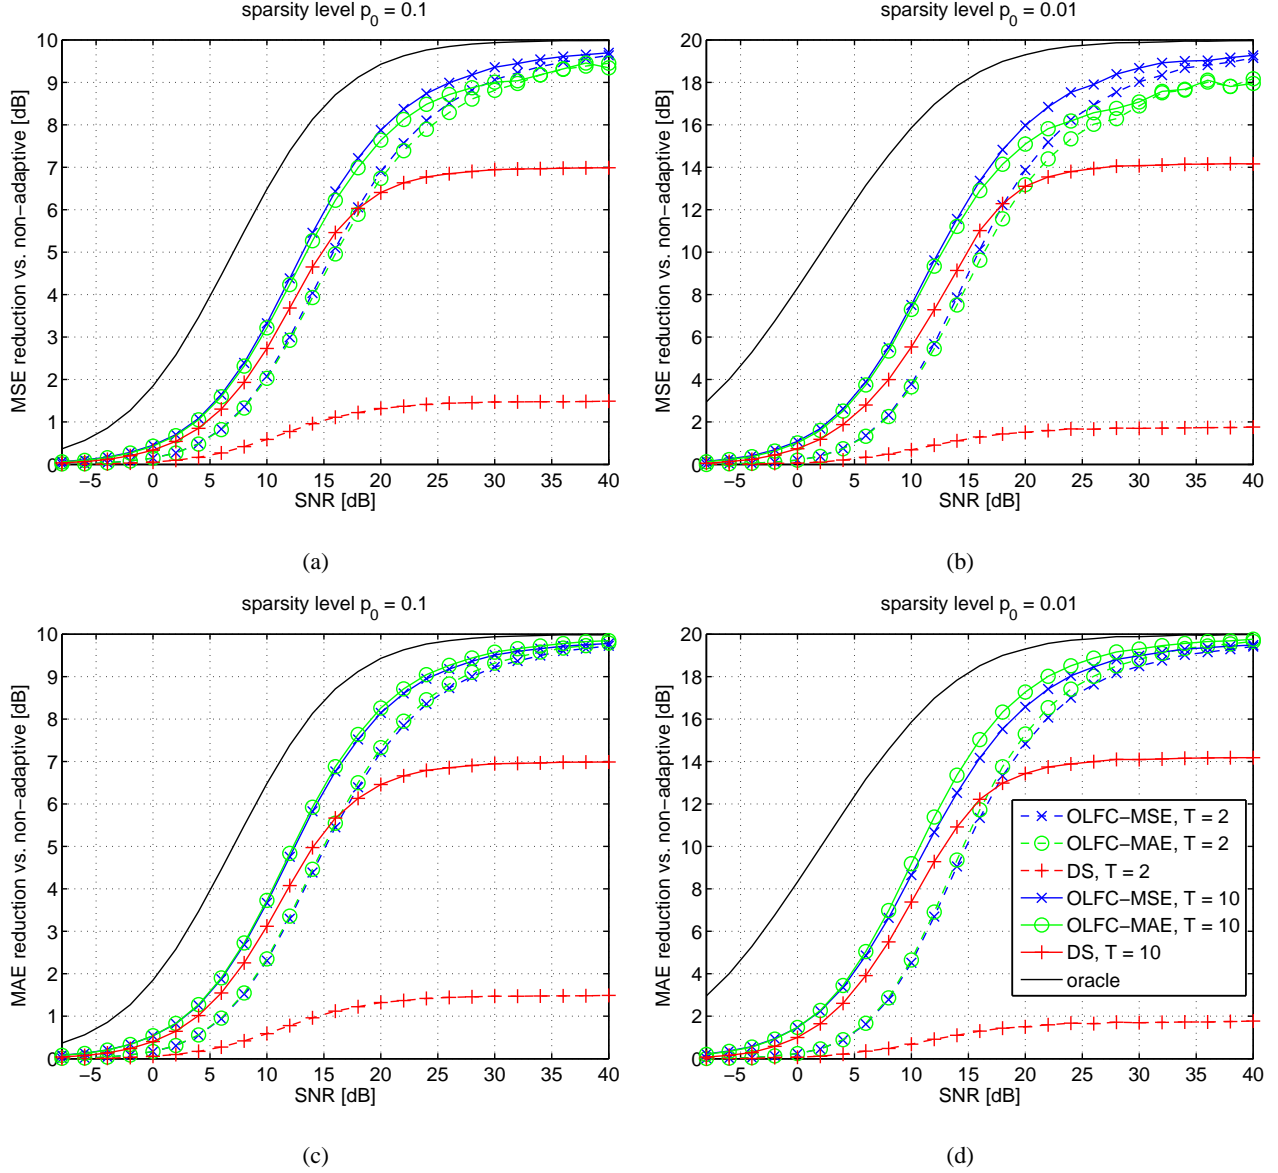


Fig. 1. Reduction in MSE (first row) and MAE (second row) relative to non-adaptive estimation as a function of SNR. The 10-stage generalized open-loop feedback control (OLFC) policies improve upon the 2-stage OLFC policies with maximum gains of 1.5 dB for $p_0 = 0.1$ and 4.5 dB for $p_0 = 0.01$. The 2-stage OLFC policies are optimal for $T = 2$. As the SNR increases, the proposed OLFC policies approach the oracle gain of $1/p_0$ and outperform distilled sensing (DS) by several dB.

policies have the best performance at all SNR.

Fig. 2 shows decreases in MSE with the number of stages T . The incremental gains predicted by Proposition 2 diminish as T increases. Using more stages is more beneficial at lower SNR and higher sparsity, whereas at higher SNR most of the signal components can be located in a single step and a




Fig. 2. MSE reduction as a function of the number of stages T . Gains diminish as T increases but less quickly at lower SNR and higher sparsity. In all cases shown, the proposed OLFC-MSE policy with 5 stages performs better than a 10-stage DS policy.

two-stage OLFC policy performs almost as well as a policy with many more stages. The gains for DS do not diminish as quickly but are lower overall, never exceeding the gain of the corresponding 5-stage generalized OLFC policy.

In Fig. 3, we consider the performance improvement due to policy rollout, as guaranteed by Proposition 3. For this experiment only, the exponent γ in (21)–(24) is fixed at $2/(q+2) = 1/2$ ($q = 2$ for MSE). The dimension N is lowered to 1000 and the results are averaged over only 1000 simulations because of the higher computational complexity of rollout. For $p_0 = 0.1$ in Fig. 3(a), no decrease in MSE is seen, whereas for $p_0 = 0.01$ in Fig. 3(b), the decrease is never more than 0.6 dB. It appears therefore that for the problem at hand, the performance gained from rollout is minimal while the computational cost of the required online simulations is much greater.

Fig. 4 depicts the fraction $\alpha^{(T)}(t)$ of the total budget allocated to each stage in a 10-stage OLFC-MSE policy for different SNR levels and $p_0 = 0.01$. The fractions $\alpha^{(T)}(t)$ are related to the fractions $\beta^{(T)}(t)$




Fig. 3. Comparison of generalized and rollout OLFC policies as a function of SNR. The improvement due to rollout is minimal.

of the remaining budget through a straightforward transformation. Three regimes may be distinguished in Fig. 4(a). At very low SNR, it is difficult to identify the signal support and the allocation is close to uniform. Between 0 and 25 dB SNR, the allocation is heavily weighted toward earlier stages. As seen in Fig. 4(b), the decrease with time is reminiscent of the geometric decay prescribed by distilled sensing. Above 25 dB SNR, the support can be determined with relatively little effort and an increasing fraction of the budget is reserved for the last stage to exploit this knowledge.

The proposed policies are based on a Bayesian framework and are thus dependent on prior knowledge of the expected sparsity level and SNR, specifically in the form of the parameters p_0 , μ_0 , and σ^2 . If these prior parameters are misspecified, the correct values can be learned through the Bayesian update process (8) but some degradation in performance is to be expected. To assess the effect of mismatched priors on the generalized OLFC policies, a series of experiments are conducted in which one of p_0 , μ_0 , or σ^2 is misspecified. In Fig. 5(a), the true sparsity level p_0 is 0.1 while the value p'_0 assumed by the policies is either 0.1 or 0.01. The performance loss of the OLFC-MSE policies is rather mild given the order-of-magnitude underestimate of p_0 . Similar results are seen when p_0 is overestimated. DS on the other hand does not make use of the parameter p'_0 and is therefore unaffected.

In Figs. 5(b) and 5(c), p_0 is set to 0.01 while the signal mean μ'_0 assumed by the policies is either correct or off by ± 4 dB. The signal standard deviation σ_0 is also changed to 0.40, making the mean mismatches on the order of one standard deviation. As can be seen from (8b), a misspecification of μ_0 leads to a

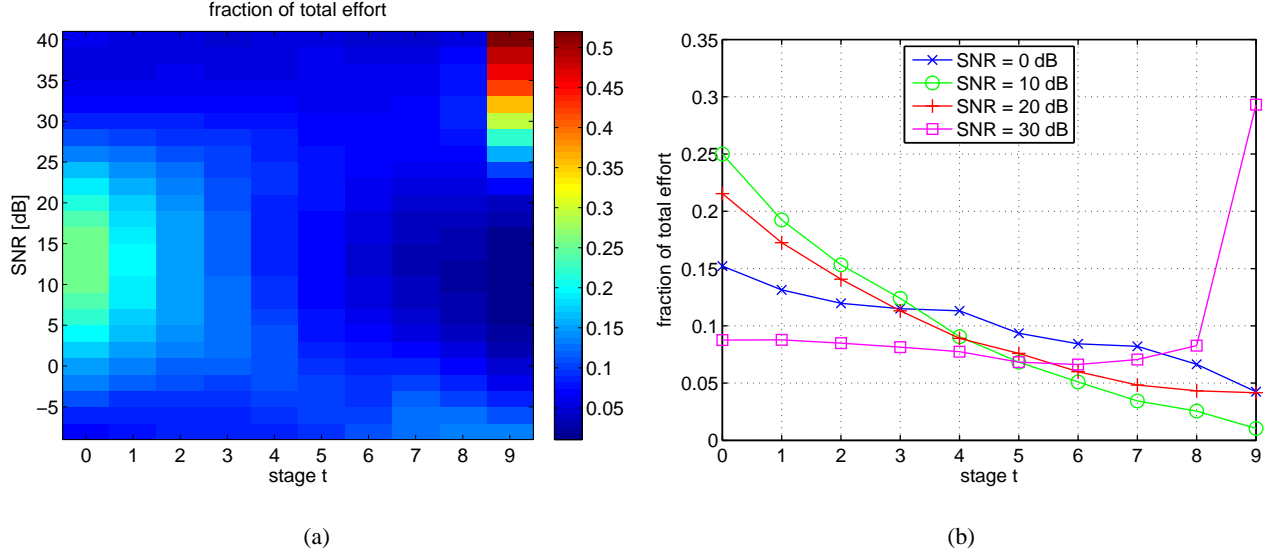


Fig. 4. Fraction of total budget allocated to each stage in a 10-stage OLFC-MSE policy for $p_0 = 0.01$ and (a) all SNR values, (b) SNR = 0, 10, 20, 30 dB. Three regimes can be seen in (a): a nearly uniform regime below 0 dB SNR, a decaying “distilling” regime between 0 and 25 dB, and a near-oracle regime above 25 dB with increasing emphasis on the last stage.

biased estimate $\mu(T)$, although the bias can be reduced by allocating more effort to the observations. It is clear from Figs. 5(b)(c) that overestimating μ_0 results in more significant losses due to missed detections of weaker than expected signal components, especially at high SNR. In contrast, when μ_0 is underestimated, the reduction in MSE relative to nonadaptive sampling can actually be greater than in the matched case; this can be attributed to a reduction in bias. In both cases, the OLFC-MSE policy remains better than DS. The consequences of misspecifying σ^2 are less severe than for μ_0 with underestimating σ^2 being worse. These findings suggest that the policies are more sensitive to overestimates of the SNR than underestimates.

V. APPLICATION TO RADAR IMAGING

In this section, the proposed allocation policies are applied to a radar imaging example also considered in [3]. The original synthetic aperture radar (SAR) image [27] shows 13 tanks in a large field and is therefore sparse in terms of targets. In the adaptive setting, it is assumed that the position and dwell time of the radar beam can be controlled, and our goal is to illustrate the benefits of such adaptivity in acquiring sparse targets.

We assume a Swerling II target model, commonly used in radar [28], in which the observation $z_i(t)$

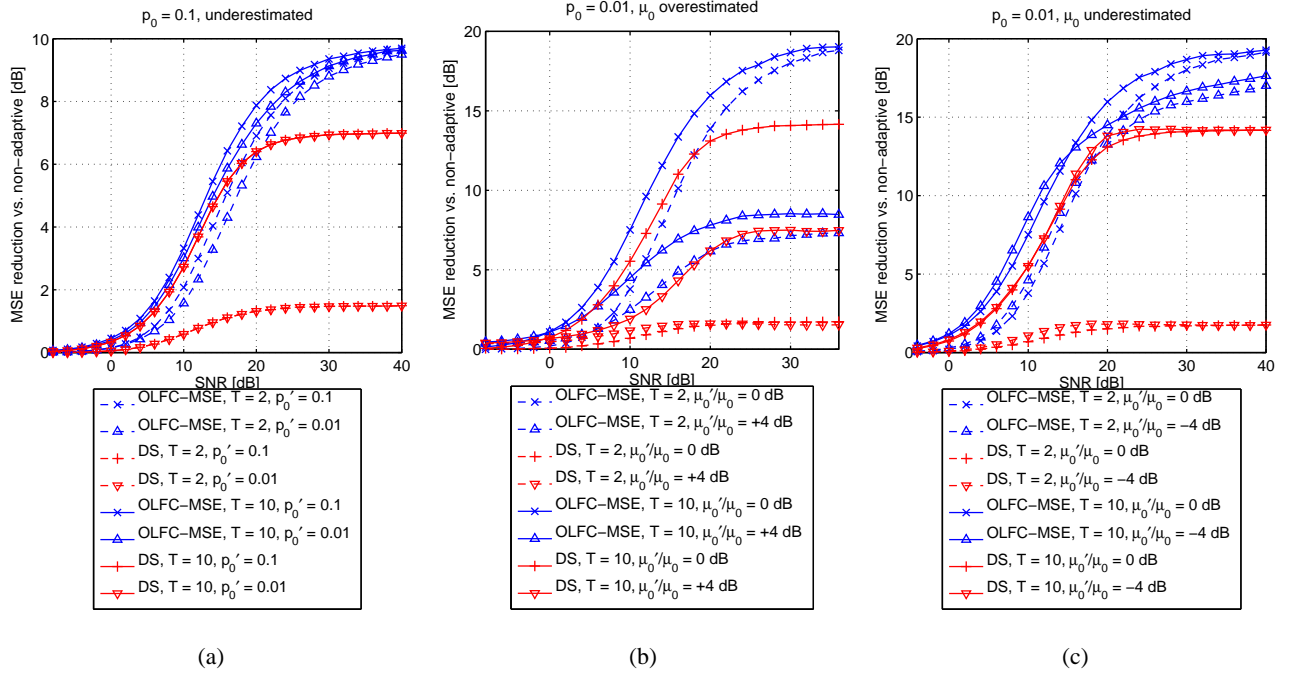


Fig. 5. Reduction in MSE relative to non-adaptive estimation as a function of SNR under mismatches in prior parameters. In (a), p_0 is underestimated by an order of magnitude and the effects are minor. More severe losses are seen in (b) with a 4 dB overestimate of μ_0 . In (c), μ_0 is underestimated by 4 dB and the losses are again modest. In all cases, the proposed OLFC-MSE policy remains better than DS.

of location i in stage t is given by the empirical mean

$$z_i(t) = \frac{1}{\kappa_i(t-1)} \sum_{s=1}^{\kappa_i(t-1)} z_{is}(t), \quad (28)$$

where the $z_{is}(t)$ are i.i.d. exponential random variables with mean equal to the true target amplitude x_i , and $\kappa_i(t-1)$ is the number of radar pulses. The total budget consists of NP pulses and the average number of pulses per location P is thus equivalent to SNR.

The Swerling observation model presents a test of robustness of the policies to non-Gaussianity. Results obtained under Gaussian and speckle noise are similar. In addition, several accommodations are made to better conform to the model in Section II. Most notably, while the targets in Fig. 6(a) are indeed sparse, they have non-negligible spatial extent and their amplitudes are not uniformly different from the background. Therefore each observed image $\mathbf{z}(t)$ is preprocessed with a matched filter, following the approach in [3] and using the same approximate tank template as in [3, Fig. 6]. The effort allocation policies are then applied with the filtered images $\mathbf{y}(t)$ as input. We use $p_0 = 0.001$ as the initial sparsity estimate in the filtered domain. The initial signal mean μ_0 and variance σ_0^2 are estimated from the first-

stage observations $y_i(1)$ above the $1 - p_0$ quantile, while the background mean (generally nonzero) and variance σ^2 are estimated from the $y_i(1)$ below the $1 - p_0$ quantile. Once the allocation $\lambda(t)$ has been determined, it is mapped to a pulse allocation $\kappa(t)$ in the unfiltered domain by convolving $\lambda(t)$ as an image with the support of the tank template (a binary image) and normalizing so that $\sum_i \kappa_i(t) = \sum_i \lambda_i(t)$. The allocation $\kappa(t)$ is then rounded to satisfy the integer restriction, again while preserving the sum.

The reconstructed image $\hat{\mathbf{x}}$ is formed as a maximum-likelihood estimate of \mathbf{x} based on $\mathbf{z}(1), \dots, \mathbf{z}(T)$:

$$\hat{x}_i = \frac{\sum_{t=1}^T \kappa_i(t) z_i(t)}{\sum_{t=1}^T \kappa_i(t)}, \quad i = 1, \dots, N.$$

In the non-adaptive single-stage case, this reduces to $\hat{x}_i = z_i(1)$ with $\kappa_i(0) = P$ in (28). Fig. 6 shows a 120×120 portion of the original image (the full 450×570 image is used in processing) together with reconstructions from $P = 3$ pulses per location. We focus attention on the targets of interest, namely the tanks. In the non-adaptive reconstruction in Fig. 6(b), the tanks are obscured by noise. Better images result from the two adaptive policies. The OLFC reconstruction however shows greater noise suppression around each tank and recovers amplitude details more faithfully.

In Fig. 7, we show one-dimensional profiles passing through the line of tanks. The middle curves indicate the true image intensities while the upper and lower curves correspond to one standard deviation above and below the mean reconstruction for each policy, where the mean and standard deviation are computed from 100 realizations. The number of pulses per location is $P = 2$. The variability in the reconstruction is clearly reduced using OLFC, in particular in the higher-amplitude regions corresponding to targets. The 5-stage OLFC policy further reduces the standard deviation by 2–3 dB relative to the 2-stage OLFC policy.

VI. CONCLUSIONS AND FUTURE WORK

We have presented multistage resource allocation policies for the sequential estimation of sparse signals under a variety of loss and effort functions. Our formulation of the problem permits the application of techniques from dynamic programming, in particular open-loop feedback control. The proposed policies improve monotonically with the number of stages and thus extend the optimal two-stage policy developed in [3]. Simulations and a radar imaging example also show gains relative to distilled sensing [10] and dramatic improvements relative to non-adaptive sensing.

The dynamic programming approach taken in this paper is quite general and can potentially be leveraged to develop tractable policies for other inference tasks such as detection or a combination of detection and estimation. More general observation models involving linear combinations may also be incorporated; the

matched filtering in the radar example in Section V is only a preliminary step in this direction. On the more theoretical side, the performance curves in Fig. 1 motivate the need for bounds on the achievable performance of adaptive sensing that are more refined than the oracle bound. Results in this vein for the case of a discrete resource budget have appeared recently [29].

APPENDIX A

DERIVATION OF POSTERIOR PROBABILITY DISTRIBUTIONS

This appendix presents the derivation of posterior probability distributions related to the observation model in (1) and the specification of a dynamic program state. Attention is paid to the adaptive nature of the observations, specifically the dependence of the sensing effort $\lambda(t)$ on past observations $\mathbf{Y}(t)$.

First we derive the conditional distribution $f(\boldsymbol{\theta} \mid \mathbf{I}, \mathbf{Y}(t))$. This can be done inductively starting with $t = 0$, in which case there are no observations and $f(\boldsymbol{\theta} \mid \mathbf{I}, \mathbf{Y}(t))$ is given by the assumed independent Gaussian prior:

$$f(\boldsymbol{\theta} \mid \mathbf{I}) = \prod_{i=1}^N f(\theta_i \mid I_i) = \prod_{i=1}^N \phi(\theta_i; \mu_i(0), \sigma_i^2(0)). \quad (29)$$

Next we assume that $f(\boldsymbol{\theta} \mid \mathbf{I}, \mathbf{Y}(t-1))$ is given and use Bayes' rule to obtain the proportionality

$$f(\boldsymbol{\theta} \mid \mathbf{I}, \mathbf{Y}(t)) \propto f(\mathbf{y}(t) \mid \boldsymbol{\theta}, \mathbf{I}, \mathbf{Y}(t-1)) f(\boldsymbol{\theta} \mid \mathbf{I}, \mathbf{Y}(t-1)) \quad (30)$$

as functions of $\boldsymbol{\theta}$. Since conditioning on $\mathbf{Y}(t-1)$ also fixes $\lambda_i(t-1)$ in (1), the observations $y_i(t)$ are conditionally independent and Gaussian and the likelihood term $f(\mathbf{y}(t) \mid \boldsymbol{\theta}, \mathbf{I}, \mathbf{Y}(t-1))$ simplifies to

$$f(\mathbf{y}(t) \mid \boldsymbol{\theta}, \mathbf{I}, \mathbf{Y}(t-1)) = \prod_{i=1}^N \phi(y_i(t); I_i \theta_i, \sigma^2/h(\lambda_i(t-1))). \quad (31)$$

From (29)–(31) it can be seen that $\boldsymbol{\theta} \mid \mathbf{I}, \mathbf{Y}(t)$ retains an independent Gaussian distribution for all t with marginals given by

$$f(\theta_i \mid I_i, \mathbf{Y}(t)) \propto \phi(y_i(t); I_i \theta_i, \sigma^2/h(\lambda_i(t-1))) f(\theta_i \mid I_i, \mathbf{Y}(t-1)). \quad (32)$$

For $I_i = 0$, θ_i is independent of $\mathbf{Y}(t)$ and $f(\theta_i \mid I_i = 0, \mathbf{Y}(t))$ is given by the i th term in (29). For $I_i = 1$, we parameterize $f(\theta_i \mid I_i = 1, \mathbf{Y}(t))$ by its mean $\mu_i(t)$ and variance $\sigma_i^2(t)$ as in Section II. A straightforward calculation based on (32) shows that $\mu_i(t)$ and $\sigma_i^2(t)$ obey the recursions in (8b) and (8c). Solving (8c) for the final-stage variance yields (6).

We now consider the conditional probability mass function $p(\mathbf{I} \mid \mathbf{Y}(t))$, proceeding inductively as before. The case $t = 0$ corresponds to the prior distribution, assumed to be i.i.d. Bernoulli:

$$p(\mathbf{I}) = \prod_{i=1}^N p(I_i) = \prod_{i=1}^N p_i(0)^{I_i} (1 - p_i(0))^{1-I_i}. \quad (33)$$

Next we relate $p(\mathbf{I} \mid \mathbf{Y}(t))$ to $p(\mathbf{I} \mid \mathbf{Y}(t-1))$ using Bayes' rule:

$$p(\mathbf{I} \mid \mathbf{Y}(t)) = \frac{f(\mathbf{y}(t) \mid \mathbf{I}, \mathbf{Y}(t-1))p(\mathbf{I} \mid \mathbf{Y}(t-1))}{\sum_{\mathbf{I}'} f(\mathbf{y}(t) \mid \mathbf{I}', \mathbf{Y}(t-1))p(\mathbf{I}' \mid \mathbf{Y}(t-1))}. \quad (34)$$

As before, conditioning on $\mathbf{Y}(t-1)$ fixes $\lambda_i(t-1)$ in (1) and thus $\mathbf{y}(t) \mid \mathbf{I}, \mathbf{Y}(t-1)$ is a linear combination of the independent Gaussian random vectors $\boldsymbol{\theta} \mid \mathbf{I}, \mathbf{Y}(t-1)$ and $\mathbf{n}(t)$. Consequently we obtain

$$f(\mathbf{y}(t) \mid \mathbf{I}, \mathbf{Y}(t-1)) = \prod_{i=1}^N \phi(y_i(t); I_i \mu_i(t-1), I_i \sigma_i^2(t-1) + \sigma^2/h(\lambda_i(t-1))), \quad (35)$$

recalling that $f(\theta_i \mid I_i = 1, \mathbf{Y}(t-1))$ is parameterized by $\mu_i(t-1)$ and $\sigma_i^2(t-1)$. From (33)–(35) it can be concluded that the components of $\mathbf{I} \mid \mathbf{Y}(t)$ remain independent with marginal distributions

$$p(I_i \mid \mathbf{Y}(t)) = \frac{f(y_i(t) \mid I_i, \mathbf{Y}(t-1))p(I_i \mid \mathbf{Y}(t-1))}{\sum_{I'_i=0}^1 f(y_i(t) \mid I'_i, \mathbf{Y}(t-1))p(I'_i \mid \mathbf{Y}(t-1))}. \quad (36)$$

The recursion for $p_i(t) = \Pr(I_i = 1 \mid \mathbf{Y}(t))$ in (8a) follows from (35) and (36).

APPENDIX B

OPTIMALITY OF THE MEAN ESTIMATOR FOR SYMMETRIC NON-DECREASING LOSS FUNCTIONS

In this appendix, we show that the expected loss in estimating a random variable θ is minimized by the mean μ of θ when the loss function is symmetric and non-decreasing and the probability density $f(\theta)$ is symmetric and strictly unimodal, as in the Gaussian case. First we prove the statement for loss functions of the form

$$L_\delta(a) = \begin{cases} 0, & 0 \leq a < \delta, \\ 1, & a > \delta \end{cases} \quad (37)$$

for $\delta > 0$. The expected loss for an estimate $\hat{\theta}$ is then

$$\mathbf{E} [L_\delta(|\hat{\theta} - \theta|)] = \int_{-\infty}^{\hat{\theta}-\delta} f(\theta) d\theta + \int_{\hat{\theta}+\delta}^{\infty} f(\theta) d\theta. \quad (38)$$

The derivative of the expected loss with respect to $\hat{\theta}$ is $f(\hat{\theta} - \delta) - f(\hat{\theta} + \delta)$ and is strictly positive for $\hat{\theta} > \mu$, strictly negative for $\hat{\theta} < \mu$, and zero only at $\hat{\theta} = \mu$ given the assumptions of symmetry and strict unimodality. Thus the estimate $\hat{\theta} = \mu$ is optimal for the loss function L_δ .

A general non-decreasing loss function L can be approximated arbitrarily closely by a sum of functions of the form in (37) in a manner reminiscent of Lebesgue integration. Given a step size $\Delta L > 0$, we construct the approximation

$$\hat{L}(a) = \Delta L \sum_{k=1}^{\infty} L_{L^{-1}(k\Delta L)}(a),$$

where $L^{-1}(k\Delta L)$ denotes the smallest value of a such that $L(a) = k\Delta L$. By the linearity of expectations, the expected value of $\hat{L} \left(\left| \hat{\theta} - \theta \right| \right)$ is a sum of functions of the form in (38). Since $\hat{\theta} = \mu$ minimizes each term in the sum individually, it also minimizes the overall sum and hence the mean estimate is optimal for \hat{L} . As $\Delta L \rightarrow 0$, \hat{L} converges to L and the statement is proven for L .

APPENDIX C

SOLUTION OF PROBLEM (20)

For notational simplicity, we write p_i , r_i , λ_i , and Λ in this appendix for the quantities $p_i(t)$, $\sigma^2/\sigma_i^2(t)$, $\bar{\lambda}_i(t)$, and $\Lambda(t)$ in (20). We also use J to denote the cost function. As noted in Section III-B, (20) is a convex minimization problem subject to a simplex constraint and therefore satisfies an optimality condition similar to (18):

$$\text{if } \lambda_i^* > 0 \text{ then } \frac{\partial J}{\partial \lambda_i}(\boldsymbol{\lambda}^*) \leq \frac{\partial J}{\partial \lambda_j}(\boldsymbol{\lambda}^*) \quad \forall j \neq i. \quad (39)$$

Condition (39) implies that the optimal solution to (20) satisfies an index rule in the sense that the non-zero components of the optimal solution correspond to the largest p_i^γ/r_i , where $\gamma = 2/(q+2)$. To prove this fact, suppose that i and j are such that $\lambda_i^* > 0$ and $\lambda_j^* = 0$ but $p_i^\gamma/r_i \leq p_j^\gamma/r_j$. Then

$$\frac{\partial J}{\partial \lambda_i} = -\frac{q}{2} \frac{p_i}{(r_i + \lambda_i^*)^{1/\gamma}} > -\frac{q}{2} \frac{p_i}{r_i^{1/\gamma}} \geq -\frac{q}{2} \frac{p_j}{r_j^{1/\gamma}} = \frac{\partial J}{\partial \lambda_j},$$

contradicting the optimality condition (39). The index rule can be stated in terms of the permutation π defined in (21), which in the notation of this appendix sorts the quantities p_i^γ/r_i in non-increasing order. Specifically, we have $\lambda_{\pi(i)}^* > 0$ for $i = 1, \dots, k$ for some integer k , $\lambda_{\pi(i)}^* = 0$ for $i = k+1, \dots, N$, and $p_{\pi(k)}^\gamma/r_{\pi(k)} > p_{\pi(k+1)}^\gamma/r_{\pi(k+1)}$ strictly.

The optimality condition (39) also implies that the partial derivatives corresponding to non-zero components of the optimal solution must be equal. Hence

$$-\frac{2}{q} \frac{\partial J}{\partial \lambda_{\pi(i)}} = \frac{p_{\pi(i)}}{(r_{\pi(i)} + \lambda_{\pi(i)}^*)^{1/\gamma}} = C^{-1/\gamma}, \quad i = 1, \dots, k, \quad (40)$$

where C is a constant to be determined. A slight rearrangement of (40) yields the expression in (23) for $i = 1, \dots, k$. The value of C in (24) is obtained by summing (23) over $i = 1, \dots, k$ and noting that $\sum_{i=1}^k \lambda_{\pi(i)}^* = \sum_{i=1}^N \lambda_i^* = \Lambda$.

It remains to determine the cutoff index k . This can be done by enforcing the condition $\lambda_{\pi(i)}^* > 0$ for $i = 1, \dots, k$ and the optimality condition (39) for $j = \pi(k+1), \dots, \pi(N)$ (corresponding to the zero-valued components). The first condition is equivalent to

$$C > \frac{r_{\pi(i)}}{p_{\pi(i)}^\gamma}, \quad i = 1, \dots, k,$$

while the second is equivalent to

$$C \leq \frac{r_{\pi(i)}}{p_{\pi(i)}^\gamma}, \quad i = k+1, \dots, N. \quad (41)$$

Given the definition of π in (21), the most stringent conditions correspond to $i = k$ and $i = k+1$, i.e.,

$$\frac{r_{\pi(k)}}{p_{\pi(k)}^\gamma} < \frac{\Lambda + \sum_{i=1}^k r_{\pi(i)}}{\sum_{i=1}^k p_{\pi(i)}^\gamma} \leq \frac{r_{\pi(k+1)}}{p_{\pi(k+1)}^\gamma}, \quad (42)$$

upon substituting (24). Solving (42) for Λ yields the condition $b(k-1) < \Lambda \leq b(k)$ using the definition of $b(k)$ in (22). If $k = N$, (41) is absent and we only have the condition $\Lambda > b(N-1)$, or equivalently we may define $b(N) = \infty$. Thus the number of non-zero components k is determined by the interval $(b(k-1), b(k)]$ to which Λ belongs. This mapping from Λ to k is well-defined if $b(k)$ is a non-decreasing function of k so that the intervals $(b(k-1), b(k)]$ are non-overlapping and span the positive real line. Indeed we have

$$\begin{aligned} b(k) &= \frac{r_{\pi(k+1)}}{p_{\pi(k+1)}^\gamma} \sum_{i=1}^k p_{\pi(i)}^\gamma - \sum_{i=1}^k r_{\pi(i)} \\ &\geq \frac{r_{\pi(k)}}{p_{\pi(k)}^\gamma} \sum_{i=1}^k p_{\pi(i)}^\gamma - \sum_{i=1}^k r_{\pi(i)} \\ &= \frac{r_{\pi(k)}}{p_{\pi(k)}^\gamma} \sum_{i=1}^{k-1} p_{\pi(i)}^\gamma - \sum_{i=1}^{k-1} r_{\pi(i)} \\ &= b(k-1), \end{aligned}$$

where the inequality is due to (21).

REFERENCES

- [1] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, pp. 489–509, Feb. 2006.
- [2] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, pp. 1289–1306, Apr. 2006.
- [3] E. Bashan, R. Raich, and A. O. Hero, "Optimal two-stage search for sparse targets using convex criteria," *IEEE Trans. Signal Process.*, vol. 56, pp. 5389–5402, Nov. 2008.
- [4] E. Bashan, G. Newstadt, and A. O. Hero, "Two-stage multiscale search for sparse targets," *IEEE Trans. Signal Process.*, vol. 59, pp. 2331–2341, May 2011.
- [5] A. Tajer, R. M. Castro, and X. Wang, "Adaptive sensing of congested spectrum bands," *IEEE Trans. Inf. Theory*, vol. 58, no. 9, pp. 6110–6125, Sep. 2012.
- [6] W. Zhang, A. K. Sadek, C. Shen, and S. J. Shellhammer, "Adaptive spectrum sensing," in *Proc. Info. Theory Appl. Workshop (ITA)*, 2010, pp. 1–7.
- [7] S. Zehetmayer, P. Bauer, and M. Posch, "Optimized multi-stage designs controlling the false discovery or the family-wise error rate," *Statist. Med.*, vol. 27, pp. 4145–4160, 2008.

- [8] G. Newstadt, E. Bashan, and A. O. Hero, "Adaptive search for sparse targets with informative priors," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, Mar. 2010, pp. 3542–3545.
- [9] D. Hitchings and D. A. Castanon, "Adaptive sensing for search with continuous actions and observations," in *Proc. IEEE Conf. Decision and Control (CDC)*, Dec. 2010, pp. 7443–7448.
- [10] J. Haupt, R. M. Castro, and R. Nowak, "Distilled sensing: Adaptive sampling for sparse detection and estimation," *IEEE Trans. Inf. Theory*, vol. 57, pp. 6222–6235, Sep. 2011.
- [11] M. Malloy and R. Nowak, "Sequential analysis in high-dimensional multiple testing and sparse recovery," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Aug. 2011, pp. 2661–2665.
- [12] —, "On the limits of sequential testing in high dimensions," in *Conf. Rec. Asilomar Conf. Signals Syst. Comput.*, Nov. 2011, pp. 1245–1249.
- [13] J. Haupt, R. Baraniuk, R. Castro, and R. Nowak, "Sequentially designed compressed sensing," in *Proc. IEEE Statist. Signal Process. Workshop (SSP)*, Aug. 2012, pp. 1–4.
- [14] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. Signal Process.*, vol. 56, pp. 2346–2356, Jun. 2008.
- [15] R. M. Castro, J. Haupt, R. Nowak, and G. M. Raz, "Finding needles in noisy haystacks," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, Apr. 2008, pp. 5133–5136.
- [16] A. Aldroubi, H. Wang, and K. Zarringhalam, "Sequential adaptive compressed sampling via Huffman codes," 2009, preprint.
- [17] M. Iwen and A. H. Tewfik, "Adaptive group testing strategies for target detection and localization in noisy environments," IMA Preprint Series, Tech. Rep. 2311, Jun. 2010.
- [18] P. Indyk, E. Price, and D. P. Woodruff, "On the power of adaptivity in sparse recovery," in *Proc. IEEE Symp. Found. Comput. Sci. (FOCS)*, Oct. 2011, pp. 1–16.
- [19] D. M. Malioutov, S. R. Sanghavi, and A. S. Willsky, "Sequential compressed sensing," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 435–444, Apr. 2010.
- [20] A. Averbuch, S. Dekel, and S. Deutsch, "Adaptive compressed image sensing using dictionaries," *SIAM J. Imaging. Sci.*, vol. 5, no. 1, pp. 57–89, 2012.
- [21] R. Castro, R. Willett, and R. Nowak, "Faster rates in regression via active learning," in *Neural Information Processing Systems (NIPS)*, 2005.
- [22] R. Willett, A. Martin, and R. Nowak, "Backcasting: Adaptive sampling for sensor networks," in *Information Processing in Sensor Networks (IPSN)*, Apr. 2004.
- [23] R. Rangarajan, R. Raich, and A. O. Hero, "Optimal sequential energy allocation for inverse problems," *IEEE J. Sel. Topics Signal Process.*, vol. 1, pp. 67–78, Jun. 2007.
- [24] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Nashua, NH: Athena Scientific, 2005, vol. 1.
- [25] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.
- [26] D. P. Palomar and J. R. Fonollosa, "Practical algorithms for a family of waterfilling solutions," *IEEE Trans. Signal Process.*, vol. 53, no. 2, pp. 686–695, Feb. 2005.
- [27] [Online]. Available: http://www.sandia.gov/RADAR/images/rtv_tanks_9in.jpg
- [28] H. Meikle, *Modern radar systems*. Artech House, 2008.
- [29] E. Arias-Castro, E. J. Candes, and M. A. Davenport, "On the fundamental limits of adaptive sensing," Nov. 2011, arXiv preprint.




Fig. 6. Portion of original image (a) in radar imaging example and reconstructions from $P = 3$ pulses per location allocated non-adaptively (b), using 5-stage DS (c), and using 5-stage OLFC-MSE (d). OLFC suppresses noise more strongly around each tank and recovers details more faithfully.

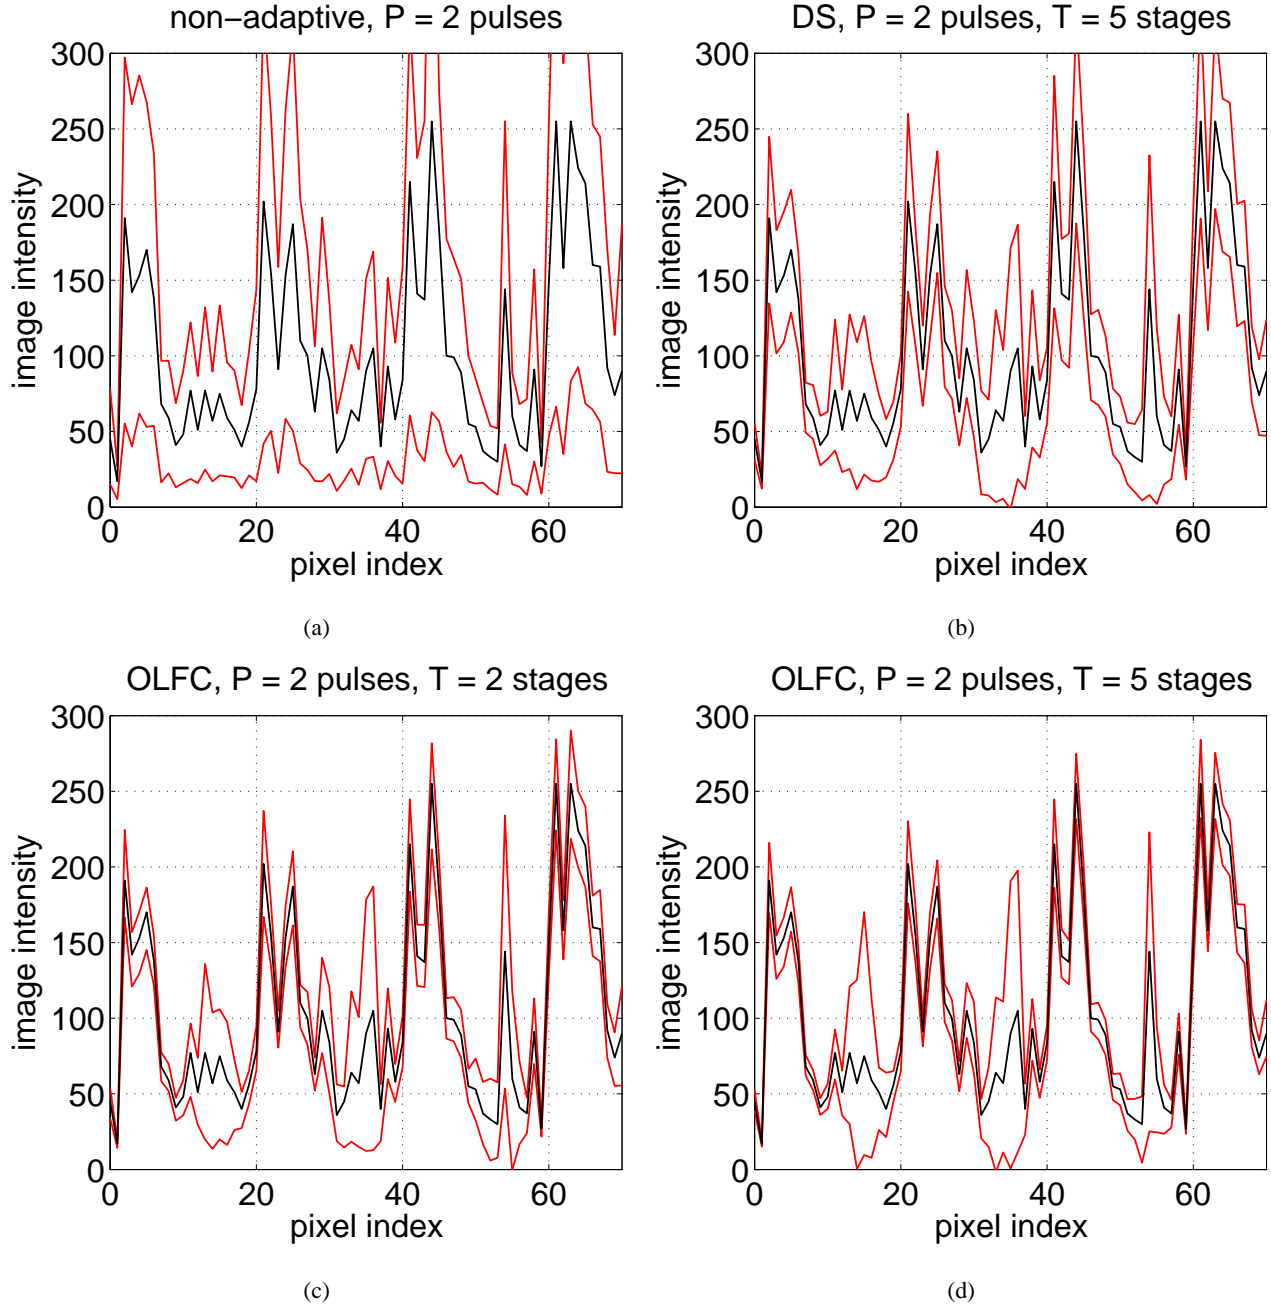


Fig. 7. One-dimensional profiles passing through the line of tanks in Fig. 6(a). Middle curves indicate the true image intensities while upper and lower curves correspond to one standard deviation above and below the mean of 100 reconstructions for each policy using $P = 2$ pulses per location. The 5-stage OLFC policy reduces the standard deviation by a further 2–3 dB relative to the 2-stage OLFC policy.